A programme to determine the exact interior of any connected digital picture

A. E. Fabris[a], V. Ramos Batista[b]

[a] *USP - IME, r. do Matão 1010, 05508-090 São Paulo-SP, Brazil*
[b] *UFABC - CMCC, r. Sta. Adélia 166, Bl.B, 09210-170 St. André-SP, Brazil*

arXiv:1401.3385v1 [cs.CG] 14 Jan 2014

Abstract

Region filling is one of the most important and fundamental operations in computer graphics and image processing. Many filling algorithms and their implementations are based on the Euclidean geometry, which are then translated into computational models moving carelessly from the continuous to the finite discrete space of the computer. The consequences of this approach is that most implementations fail when tested for challenging degenerate and nearly degenerate regions. We present a correct integer-only procedure that works for all connected digital pictures. It finds all possible interior points, which are then displayed and stored in a locating matrix. Namely, we present a filling *and* locating procedure that can be used in computer graphics and image processing applications.

Keywords: Raster Graphics, Region Filling, Point-based Modelling, Discrete Jordan Curve Theorem

1. Introduction

Region filling is a crucial operation in computer graphics and image processing. There are two basic approaches to region filling on raster systems. One of them is to start from a given interior position called *seed*, and paint outward from this point until we meet the specified boundary conditions. The other is to determine the overlap intervals of scan lines that cross the region. The seed-fill methods are mostly applied in interactive painting systems. Alternatively, the scan-conversion approach is typically used in hardware implementations and in general graphic packages to fill polygons, circumferences, ellipses and spline curves. These two approaches will be resumed in Section 2, but now we focus on their limitations: the non-robustness of scan conversion, and the high computational cost of seed fill.

As explained in Section 2, neither seed fill nor scan conversion can really provide perfect filling at the lowest computational cost. This is yet desirable in many circumstances. We list three of them.

The first is represented in Figure 1. It schematises the mould of an integrated circuit. In general, thousands of models are manufactured from the mould. A single flaw in it will produce useless lots, and therefore a great waste of material and time. Until nowadays, this kind of flaw detection still depends on handwork.

In fact, Figure 1 right is an output of our algorithms. They aim at filling regions perfectly and quickly. As an application, this helps detect imperceptible flaws. For the time being, however, they can detect only *external* flaws, but are adaptable for the whole job.

Our second example is of theoretical interest. Of course, there is already a valuable and extensive literature about region filling and point location. In fact, both problems are intimately related. Now, according to [1], none of them is free of minor failures, and the best ones are computationally expensive. Therefore, if there are algorithms that prove to be quick and faultless, they are really worth studying.

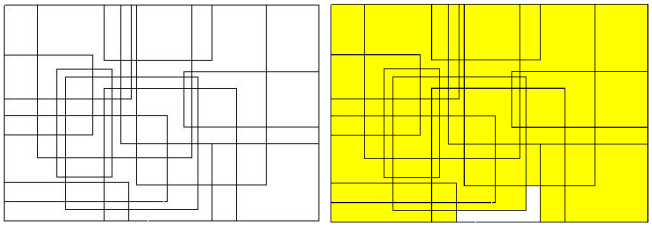

Figure 1: Detecting an imperceptible flaw in a mould of integrated circuit.

The third and most striking example is the technology of 3D-printing. This paper deals with a two-dimensional approach. But our algorithms were elaborated with ideas that *are* extendable to higher dimensions. In fact, we are already researching this problem for future works. The basic idea is to slice a 3-dimensional object into several 2-dimensional ones, which are then treated individually. In this case, least computational time becomes highly important. But also perfect filling and point location. The problem of printing plastic cups, for instance, is equivalent to what Figure 1 illustrates.

Of course, when performing region filling or point location we can content ourselves with flaws as we see in Figure 2 drawn in Xfig and saved as a PostScript file. If not, seed fill reduces them to a minimum but usually with a high computational cost.

Our research was motivated by the fact that scan-conversion algorithms are not always reliable. In the case of point location algorithms, they fail even for non-challenging regions, as demonstrated in [1, 2]. Moreover, in [1] Schirra presents a complete comparison among all classical point-location algorithms.

In this paper, the difficulties in discretisation are acknowl-

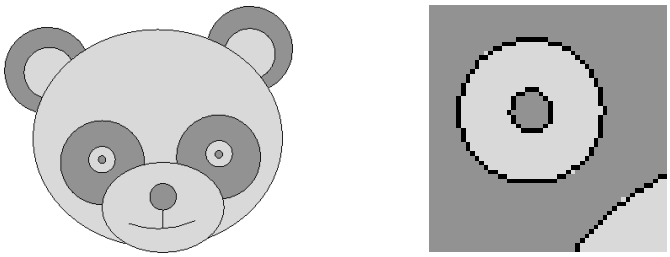

Figure 2: (a) A panda's drawing; (b) detail of area around one eye showing individual pixels.

edged and tackled by casting the problem to be solved as a discrete integer problem from the very outset. We believe that region filling and point location algorithms cannot be always reliable, unless these problems are set up and solved in a discrete space. Hence, we *never* have to accommodate any inconvenience of numerical accuracy that follows from floating point discretisation, geometric approximation, or less rigorous approach to final discretisation in the rendering process. We give formal definitions of the involved discrete topology, together with a mathematical proof of robustness.

However, we *do not* start from mathematical models. Our input is solely a matrix with entries 0 and 1, which we interpret as a picture of black pixels surrounded by white pixels. No equations are given beforehand. Then we *construct* a mathematical model that retrieves our input, performs the region filling and generates a filling matrix. This matrix locates any point as belonging to the interior, exterior or to the picture itself.

Our work is organised as follows. Section 2 gives some background on the two classical region filling algorithms. Section 3 explains some basic terms and definitions. They are used in Section 4, where we discuss our main programme and its algorithm. The main programme is endowed with an optional improvement, which is explained in Section 5. This improvement will make the main programme correct, and Section 6 presents a theoretical proof of this fact. Details about downloading and running our programme are explained in Section 7, whereas Section 8 is devoted to our conclusions. Finally, Appendix A describes some theoretical results that can be achieved through the source codes.

2. Seed fill and scan conversion

Seed-fill algorithms can be subdivided into four steps [3]: a *start procedure* to choose an interior starting point, namely the seed; a *propagation method* to determine the next considered point; an *inside test* to check whether a point is in the region and if it should be marked; a *set procedure* to change the colour of a given picture element or *pixel*. The *start procedure* can be done automatically or user-guided. For complex objects, multiple seeds may be required and their automatic detection can be very difficult. This is one of the main reasons why this family of algorithms is more suitable for interactive region filling [4]. For the *propagation method*, the naïve implementation of the algorithm searches for connected pixels recursively, and can be based either on pixel- or line-adjacency. In the former, the filling procedure is to move from a current pixel to all immediate neighbours in a certain order. In the latter, the interior of a region is regarded as a set of adjacent horizontal line segments. Usually, seed-fill algorithms need working memory in addition to the frame buffer, and are slower than scan-conversion algorithms. The amount of working memory cannot be constant, for it increases with both size and complexity of the region. Consequently, it cannot be stored in advance. Compared to seed-fill algorithms, scan-conversion algorithms are generally faster and require less or no additional working memory. In order to overcome the difficulties of seed-fill algorithms, different strategies like [5, 6] try to re-model the region and so get a fast image manipulation, processing and displaying.

Regarding scan conversion, well-known related to algorithms require computationally expensive presorting and marking phases before they compute the actual intervals of points contained in the region. This holds even for simple concave polygons. For curved boundaries, the marking phase requires careful attention to both geometric and numerical details to provide robust algorithms. Conventional filling of curved regions is based on algorithms for scan conversion of polygons, in which the end points of spans are incrementally updated, and then the intermediate pixels are filled. At some stage the intersection formula, derived in the continuous plane and usually computed using real arithmetic, has to be mapped to the discrete plane. This mapping needs an implicit or explicit epsilon-test that may cause incorrect results. Many scan-conversion algorithms fail to handle complex objects correctly, while in principle seed filling is more robust [1, 4].

Levoy and Whitted suggest points rather than triangle meshes as display primitives [7, 8, 9]. Similarly, we propose points as low-level modelling primitives, and so get rid of the well-known discretisation problems of the polygon and of all general curved boundaries too. Points are building bricks of geometry discretisation, and they can also provide a common format at the lowest level for both constructed and scanned geometric data. Namely, for both computer graphics and image processing.

3. The mathematical structure

A dominant reason for which raster graphics algorithms fail in robustness is the absence of formal specifications. Raster systems are discrete devices, which generate image by displaying the intensity value of each pixel in a finite two-dimensional matrix of pixels. But discrete definitions are implicitly assumed to be like their respective counterparts of the continuous Euclidean space. Commonly, a formal definition of interior/exterior of digital objects is omitted. This can cause difficulties in differing between the rules of odd-parity and nonzero winding number, and mainly in comparing different algorithms. Figure 3 shows a typical ambiguity that arises when such terms are treated only implicitly.

Now we define some basic terms, and the first is our space of points:

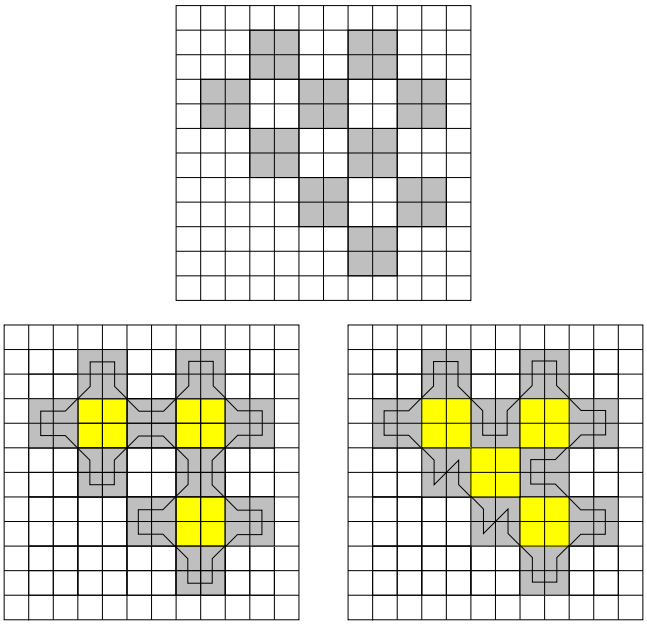

Figure 3: Ambiguity regarding interior and exterior of a digital picture (example with a 12×12 matrix of pixels).

Definition 3.1. Given any two positive natural numbers N and M, a *canvas* in $\mathbb{Z}^2$ is the set $\{(p_y, p_x) : 1 \leq y \leq N, 1 \leq x \leq M\}$, where $N \times M$ indicates its *dimension*.

In a computing context, a digital image is a two-dimensional array of values reflecting the brightness and colours of the pixels.

Definition 3.2. Let $S = S_{N\times M}$ be a canvas. For any $f : S \to \{0, 1\}$, we say that the set $f^{-1}(1)$ is a *digital picture* in the canvas.

In Definition 3.2, our convention is to think of 0 as white and 1 as black. Therefore, the *digital picture* is a set of black dots over a white background. As remarked in the introduction, our input is a digital picture $f^{-1}(1)$, or just a matrix with entries 0 and 1. Our filling algorithm does not require any other mathematical expression for the function f.

Definition 3.3. A *discrete curve* in a canvas S is an 8-connected list $L : [0 \dots n] \to S$. We say that L is *monotone and closed* when $L(i) = L(j) \iff i = j \bmod n$.

Definition 3.4. Consider a digital picture $\mathcal{P}$ with its corresponding $f : S \to \{0, 1\}$. If any two points $p, q \in f^{-1}(1)$ always admit a discrete curve $L : [0 \dots n] \to S$ with $L(0) = p$, $L(n) = q$ and $f(L(i)) = 1\ \forall i$, then we say that $\mathcal{P}$ is *connected*.

These definitions are comparable to those given in works by Rosenfeld and Kong [10, 11]. Notice that the same digital picture can correspond to distinct discrete curves. By looking at Figure 3, the black pixels on either case form exactly the *same* digital picture $\mathcal{P}$. However, we can *track* these pixels in different ways. Each image comes from a different discrete curve that represents $\mathcal{P}$.

Figure 4 helps understand the following definition:

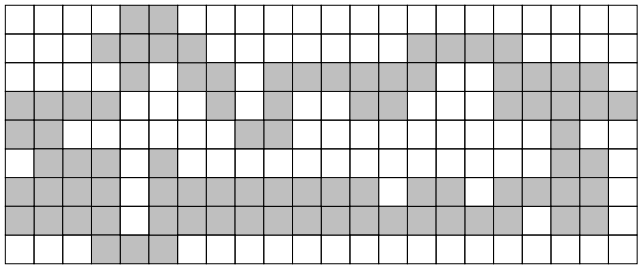

Figure 4: A digital picture with thickness varying from 1 to 3.

Definition 3.5. Given $p \in L([0 \dots n])$, consider the biggest square Q such that $p \in Q \subset L([0 \dots n])$. Its dimension is $k \times k$ for a certain $k \geq 1$. Then k is the *thickness of the discrete curve L at p*.

Definition 3.6. We say that a discrete curve L is *thin* when it has thickness 1 everywhere.

In practice, Definition 3.6 can be replaced by a weaker condition:

Definition 3.7. An injective discrete curve $L : [0 \dots n] \to S$ is *locally thin* when the following holds. There are intervals $I_1 \cup \dots \cup I_k = [0 \dots n]$ such that $|I_j| \geq 4$ and $L|_{I_j}$ is thin, $\forall\, j \in \{1, \dots, k\}$.

For instance, the black pixels in Figure 4 cannot represent the image of an L that is locally thin. We shall give more details in Section 5.

Our purpose is to consider *any* digital picture $\mathcal{P}$, solely restricted to Definition 3.4. For instance, take a more general example like in Figure 4. Now the question is: How to describe our input, a digital picture $\mathcal{P}$ as in Figure 4, in terms of a function L in accordance with Definition 3.3?

This question will be answered in Section 5. In fact, that section introduces what we call *the Lego curve*. Its set of pixels will not always coincide with $\mathcal{P}$. However, starting from the Lego curve it is then possible to obtain a function $L : [0 \dots n] \to S$ as in Definition 3.3, such that $L([0 \dots n]) \equiv \mathcal{P}$.

Notice that Definition 3.3 *does not* require a discrete curve to be *simple*. Namely, it can have *self-intersections* defined as follows:

Definition 3.8. Based on Definition 3.3, we say that L has
- a *self-crossing* when there are two disjoint intervals $I, J \subset [0, n]$, $|I| = |J| = 2$, such that $L(I) \neq L(J)$ and all pixels $L(I \cup J)$ are 8-connected.
- an *overlapping* when we have $I, J \subset [0, n]$, $\min\{|I|, |J|\} \geq 2$, such that $I \cap J$ is unitary, and either $L(I) \subset L(J)$ or $L(J) \subset L(I)$.

If any of these two cases occur, then L has a *self-intersection*. Otherwise, it is a *simple* curve.

Examples in $\mathbb{R}^2$ can clarify Definition 3.8. For instance, the so-called *Bernoulli's Lemniscate* is a closed curve, of which

the parametric equations are $x(t) = 2\cos t/(1+\sin^2 t)$, $y(t) = \sin 2t/(1+\sin^2 t)$, with $t \in [-\pi, \pi]$. Its trace is ∞ (the symbol of infinity). In this case we just have a self-crossing at the origin. For the curve $(0, \sin t)$, $t \in [0, \pi]$, overlapping occurs at $t = 0$ and at $t = \pi/2$. Notice that t could vary in a bigger domain, but we always take the *smallest* for which the curve is closed.

The reader must have noticed that we treat some technical terms as synonyms: point and pixel, discrete and digital, etc. In fact, this started to happen right at Definition 3.1, because our constructions are devoted to both theoretical and practical contexts. The reader is free to choose any of these contexts, and then think in accordance with the corresponding terms.

4. The main programme

In this section we discuss the algorithm of our main programme, named `loci.m` after a mnemonic to "location of curve interior". In Section 7 the reader will find details about downloading and running it. The programme `loci.m` is an implementation of the "Filling Up Algorithm" (FUA). Before discussing FUA, it is important to explain some hypotheses under which we claim that `loci.m` will always give a correct answer.

We recall the difference between two similar terms: *connectedness* and *connectivity*. Connectedness of a digital picture $\mathcal{P}$ is a global property implied by Definition 3.4. Connectivity is a local property that applies to each pixel with respect to the ones that surround it in a digital picture $\mathcal{P}$.

Namely, given any $p \in \mathcal{P}$, if either $p \pm (1, 0)$ or $p \pm (0, 1)$ is again in $\mathcal{P}$, then it has connectivity 4 with p. If any of $p + (\pm 1, \pm 1)$ is in $\mathcal{P}$, then it has connectivity D with p. Finally, when both 4- and D-cases apply to p, we say in short that surrounding pixels have connectivity 8 with p.

Now we summarise the *three* conditions on L that will make FUA correct. The third condition will be explained in the sequel:

C1. L is locally thin;

C2. L can only be tracked horizontally and vertically. In other words, $L : [0 \ldots n] \to S$ is such that $L(i)$ and $L(i+1)$ are 4-connected $\forall\, i < n$;

C3. L is "spike-free".

Before discussing **C3**, we need to introduce a concept called $i - o$ *pixels* (see Definition 5.1) which is analogous regarding scan-conversion algorithms to the usual marking phase required to compute the intervals of points contained in the region.

From upside down, the scanning of a digital picture gives an $N \times M$ matrix of black and white pixels. Typical rows are represented in Figure 5. We track each row from the left to the right. In this process, i-pixels indicate the entry to and o-pixels the exit from $\mathcal{P}$. When a pixel is simultaneously i and o, we mark it with an x.

In Section 5 we are going to introduce CoTRA, Connectivity and Thickness Reduction Algorithm. CoTRA produces our "Lego curve" L.

Figure 6 gives an example of the input (top) and the output (bottom) by CoTRA applied to a digital picture $\mathcal{P}$. In fact, $\mathcal{P}$

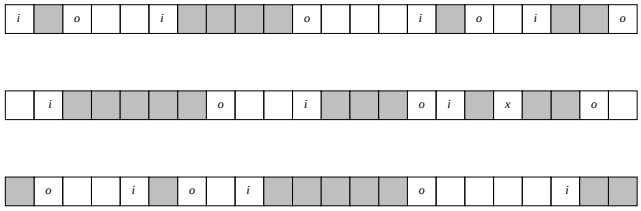

Figure 5: Some typical rows of a digital picture and their entry-exit pixels.

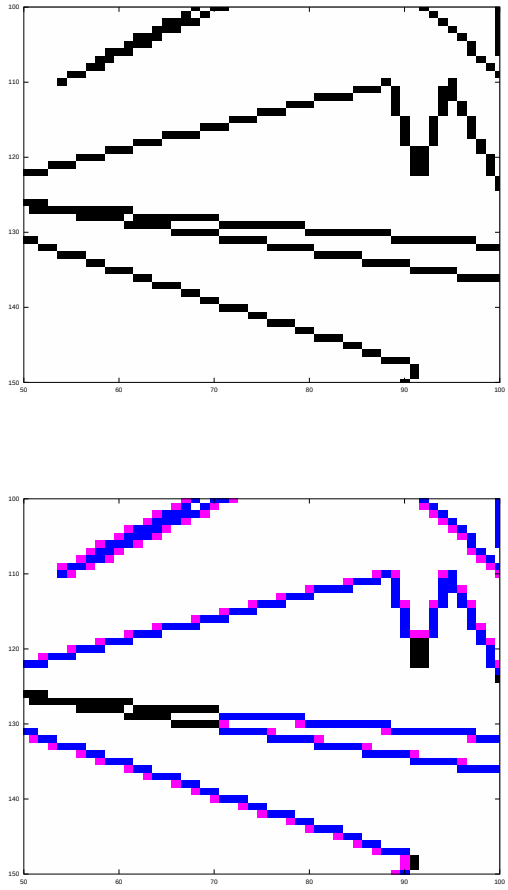

Figure 6: Input (top) and output (bottom) of the Connectivity and Thickness Reduction Algorithm (CoTRA).

is a detail of Figure 9 obtained from the drawing of a non-self-intersecting polygon in Xfig. The spurious self-intersection at the left middle occurs due to discretization errors when moving from the ideal mathematical model to the digital picture.

On top, $\mathcal{P}$ is given as a set of black pixels in the white canvas. At the bottom, we mark blue pixels that are originally from $\mathcal{P}$, and $i - o$ pixels in magenta. The blue and magenta pixels form the sought after Lego curve L. However, notice that not all pixels were marked by CoTRA. For example, close to $(y, x) = (130, 70)$ it found a stretch interpreted as a self-intersection of $\mathcal{P}$. Moreover, close to $(120, 90)$ and $(150, 90)$ we have turning points. CoTRA cuts them by keeping local thickness 1. We assert that FUA will be correct for Lego curves

L (see Section 6 for details).

In Section 5, we formalise "interior" of $\mathcal{P}$ by means of Definition 5.3.

Now the reader may be curious about the turning points of Figure 6. Of course, the one close to $(120, 90)$ could also have thickness 1 and look like a thinner "spike". In technical terms, we have

Definition 4.1. Let $L : [0 \dots n] \to S$ be a monotone discrete curve with $L(i) = (L_y(i), L_x(i)) \, \forall \, i$. For all $-3 < k < 3$, if there is i such that $|L_x(i+k) - L_x(i)| < 1$ and either $L_y(i) \geq L_y(i+k)$ or $L_y(i) \leq L_y(i+k)$, then L has a *vertical spike* at i. If there is i such that $|L_y(i+k) - L_y(i)| < 1$ and either $L_x(i) \geq L_x(i+k)$ or $L_x(i) \leq L_x(i+k)$, then L has a *horizontal spike* at i.

Handling a picture with spikes proved to needlessly increase the complexity of FUA. Indeed, we are going to obtain the interior of $\mathcal{P}$ as the interior of L *minus* the set $\mathcal{P}$ itself. CoTRA was programmed in order to eliminate spikes and then correct FUA.

As in any improvement, CoTRA adds a computational cost. But the users can choose: if the first output is not satisfactory, they call CoTRA for a second answer.

Now we present the basic algorithm to fill up $\mathcal{P}$ by marking its entry-pixels (see Algorithm 1).

Algorithm 1 The "Filling Up Algorithm" (FUA)

Input: Digital Picture $\mathcal{P} = f^{-1}(1) \subset img(:, :, 1)$
Output: Interior of $\mathcal{P}$

1: **for** $i = 2$:length$(img(:, 1, 1)) - 1$ **do**
2: $\quad bps \leftarrow$ Find_ioPixels
3: $\quad$ **for** $k = 1 : 2$:length$(bps) - 1$ **do**
4: $\quad\quad$ **if** $img(i, bps(k) : bps(k+1), 1)$ is an extreme **then**
5: $\quad\quad\quad$ shorten bps
6: **for** $k = 2 : 4$:length$(bps) - 1$ **do**
7: $\quad img(i, bps(k) : bps(k+1), 3) = 0$

Notice that the "spike-free" condition is important in lines 4-5, where we "jump" the $i - o$ pixels of a local maximum (or minimum) stretch of the picture. This is because the status of being inside/outside the picture does not change at passing through such stretches.

For the rest, lines 6-7 simply mark the *actual* interior points of the ith row in yellow (R = 255, G = 255, B = 0). In the next section, we shall formalise several terms already used herein.

5. Definition of the Lego curve

At any row of pixels, the digital picture has white and black ones. Fix any row and consider the pixels in it. Up to scaling, each of then has an x-coordinate that ranges from a natural number to its successor.

Definition 5.1. Given a row R of black and white pixels, we say that $p \in R$ is an *i-pixel* when p is white and $p + 1$ is black. We say that $p \in R$ is an *o-pixel* when p is white and $p - 1$ is black. If p is simultaneously i and o, we call it an *x-pixel*.

The letters i and o stand for the process of getting inside/outside the black horizontal segments from Figure 4, as depicted in Figure 5. Definition 5.1 implies that the i-o sequence is always alternated, if not empty. By adding a frame of white pixels around the picture, we guarantee the sequence to begin with i and end with o. We have added this frame in Figure 7.

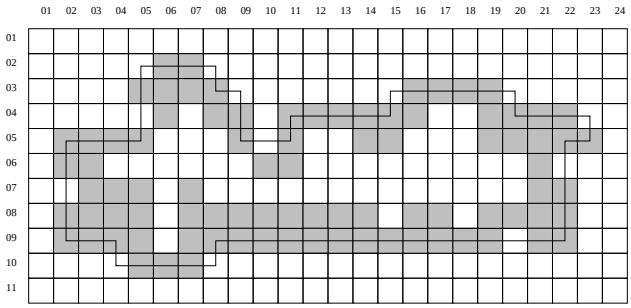

Figure 7: A digital picture and its corresponding track of the Lego curve.

We recall Definition 3.7. In Figure 7, the black pixels form a picture $\mathcal{P}$. Notice that $\mathcal{P}$ cannot be the image of a discrete curve that is locally thin. This is only possible if we remove at least 9 pixels from $\mathcal{P}$, namely $(7, 7)$ and $(8, 7 - 14)$.

In order to define the Lego curve, we first need to introduce two concepts as follows:

Definition 5.2. Consider a pixel canvas S with dimension $N \times M$, and also a digital picture $\mathcal{P}$ contained therein. Suppose $\mathcal{P}$ is surrounded by the white pixel frame $\mathcal{F} := \{1, N\} \times [1, M] \cup [1, N] \times \{1, M\}$. In this case, the *exterior* of $\mathcal{P}$ is the set $\mathcal{E}$, $\mathcal{F} \subset \mathcal{E} \subset (S \setminus \mathcal{P})$, such that for any two points $p, q \in \mathcal{E}$ there exists a discrete curve $L : [0 \dots n] \to (S \setminus \mathcal{P})$ with $L(i)$ and $L(i+1)$ 4-connected $\forall \, i < n$.

Definition 5.3. Under the hypotheses of Definition 5.2, if $\mathcal{E}$ is the exterior of $\mathcal{P}$, then its *interior* is $S \setminus (\mathcal{P} \cup \mathcal{E})$.

Definition 5.4. Let $\mathcal{P}$ be a digital picture and p a corresponding i or o-pixel, such that p belongs to the exterior of $\mathcal{P}$. If $\mathcal{P}$ has two pixels that are D-connected to each other, and also 4-connected to p, then we say that p is an *L-pixel* of $\mathcal{P}$.

In Definition 5.4, if we write the pixel coordinates as (line, column) in Figure 7, then $\mathcal{P}$ has exactly the following L-pixels: $(2, 5)$, $(4, 5)$, $(7, 2)$, $(10, 4)$, $(10, 8)$, $(9, 20)$ $(6, 22)$, $(4, 23)$, $(3, 20)$, $(3, 15)$, $(5, 10)$, $(3, 9)$ and $(2, 8)$.

Definition 5.5. Given a digital picture $\mathcal{P}$ with exterior $\mathcal{E}$, let $\mathcal{L}$ be the set of L-pixels of $\mathcal{P}$ that are 4-connected with $\mathcal{E}$. The *Lego curve* of $\mathcal{P}$ is a digital curve of minimal length that satisfies **C2**, and has an exterior $\mathcal{E}' \subset \mathcal{E}$ such that $\mathcal{E} \setminus \mathcal{E}' \equiv \mathcal{L}$.

We conclude this section by remarking that Definition 5.5 *does not* require the Lego curve to be locally thin, not even free of self-intersections. However, they can only occur in the special form of *overlapping*, not *self-crossing*. The FUA does work in the former case, but can fail in the latter.

6. Existence and uniqueness of the Lego curve

In this section we present our main theorem, which translates to the open source code `lego.m`. We have written our proof like a *pseudocode*.

Theorem 6.1. *Any digital picture $\mathcal{P}$ admits a unique corresponding Lego curve L.*

Proof: We are going to construct a list $L : [0 \dots n] \to \mathbb{Z}^2$. Take $\mathcal{P}$ as in Definition 5.2. Its pixels will be called *black*. However, we shall say that a pixel is *white* exactly when it belongs to the exterior of $\mathcal{P}$. In what follows, both the words *increase/decrease* mean "by one".

Consider the *i-o* pixels of $\mathcal{P}$ with least ordinate y. Among them, name p the one with the least abscissa x. It has coordinates (p_y, p_x), $1 < p_y < N$ and $1 < p_x < M$. We start with $n = 0$ and either $L(0) = p$ or $L(0) = p + (0, 1)$, depending on whether p is an L-pixel or not, respectively. At any step, the construction interrupts when $L(n)$ coincides with $L(0)$.

Step I. Increase n and define $L(n)$ as follows:

(a) if at the southwest of $L(n-1)$ we have a black or an L-pixel, then it will be $L(n+1)$, and $L(n)$ will be the southern pixel;
(b) otherwise, if the south is black or an L-pixel then it will be $L(n)$;
(c) if neither (a) nor (b) then take $L(n) = L(n-1) + (0, 1)$.

Step II. Now $L(n)$ has just been defined.

If $L(n+1)$ is already defined, interchange the roles $x \leftrightarrow y$ by rotating the canvas counterclockwise. Increase n, redefine p as $L(n)$ and repeat Steps I-II.

If $L(n+1)$ is not defined yet, then we look at $L(n)_y$. In the case it is $L(n-1)_y + 1$, redefine p as $L(n)$ and repeat Steps I-II. In the case it is $L(n-1)_y$, we can have

(a) $L(n)$ and $L(n) + (1, 0)$ both white; this situation is predicted in Step III.
(b) $L(n) + (1, 0)$ is black or an L-pixel; take $L(n+1) = L(n) + (1, 0)$, increase n, redefine $p = L(n)$ and repeat Steps I-II.
(c) Otherwise redefine $p = L(n)$ and repeat Steps I-II.

Step III. In the present orientation you cannot increase $L(n)_y$ any more. Then interchange the roles $x \leftrightarrow y$ by rotating the canvas clockwise. Clear $L(n)$, decrease n and repeat Steps I-II. Note: if this step happens three times consecutively, then $\mathcal{P}$ consists of a single pixel. This case can be excluded beforehand.

Remark: If $\mathcal{P}$ does not consist of a single row of black pixels, then there will be a least value $n_1 \leq n$ such that $L(n_1)$ has a greater ordinate than $L(0)$ with respect to the initial orientation.

The above construction is in fact what we called CoTRA. But the reader will notice a gap in it, as illustrated in Figure 8. It is easily identified by "trapped" 4-connected L-pixels. Compare it with Figure 6. But we easily go round this gap by identifying non-marked L-pixels and locally repeating the above construction.

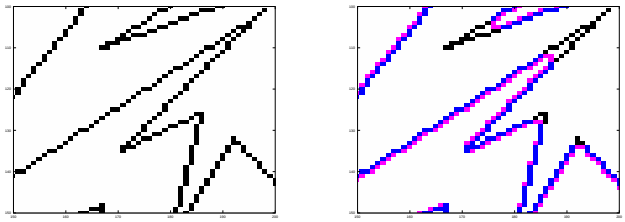

Figure 8: False self-intersections interpreted by CoTRA.

However, our list L may still have spikes. In this case, we recall Section 4 and see that **C2** trivially holds for L. Regarding **C1**, it easily follows from **C2** and **C3**. In order to *get* **C3**, we have to adapt the above construction as follows: if the roles $x \leftrightarrow y$ are interchanged consecutively twice, then lop off some newly defined $L(n)$ by checking if either x- or y-coordinates remained constant.

Of course, there remain two important questions that were not answered yet. The first is: why does the construction ever finish? In other words, why does $L(n)$ coincide with $L(0)$ for a certain positive n?

Notice that any two 8-connected pixels of $\mathcal{P}$, that are both connected to $\mathcal{E}$, can either project onto Ox or Oy injectively. We start with Oy and keep on using it until injectivity fails. Then we change to Ox and do the same to this new axis. This is exactly when the roles $x \leftrightarrow y$ are interchanged in the above construction. It begins with a local graph $x = x(y)$ so that pixels right *under* the graph belong to the exterior $\mathcal{E}$ of $\mathcal{P}$. As a matter of fact, this reasoning holds for the *whole* construction. Thus no pixel of $\mathcal{E}$, that is right under a local graph, is left out. Hence $L(n)$ will eventually coincide with $L(0)$ for a certain n.

By the way, the same reasoning explains why L contains the set $\mathcal{L}$ from Definition 5.5. This partially answers the second question: why L fits Definition 5.5? It remains to check that it has minimal length among any other curve L' with exterior $\mathcal{E}'' \subset \mathcal{E}$ and $\mathcal{E} \setminus \mathcal{E}'' = \mathcal{L}$.

We know $\mathcal{E} = \mathcal{E}' \dot\cup \mathcal{L} = \mathcal{E}'' \dot\cup \mathcal{L}$, thus $\mathcal{E}' = \mathcal{E}''$. Suppose there exists $p \in \mathcal{L} \setminus \mathcal{L}'$. Consequently, $p \in \mathcal{P}$ and is 4-connected with $\mathcal{E}$. Indeed, it cannot be only D-connected because then an L-pixel would have replaced it in the above construction. Either p is 4-connected with $\mathcal{E}' = \mathcal{E}''$ or p is simultaneously 4-connected with two different pixels $q, r \in \mathcal{L}$. This latter is the only possibility for $p \notin \mathcal{L}'$, since $p \notin \mathcal{L}$ and both L, L' have to satisfy **C2**. But then $q, r \in \mathcal{L}'$ have to be connected by a longer path than through p. Hence L' cannot have minimal length. q.e.d.

We finish this section by stating and proving a corollary, which is a weaker version of the *discrete Jordan curve theorem*. This is because we use an additional hypothesis H, described below. In Appendix A we discuss some other theoretical results that can be derived from ours.

Corollary 6.1. *Let $C : [0 \dots n] \to S$ be a simple closed curve and denote its image by $\mathcal{P}$. Suppose S and $\mathcal{P}$ are such that Definition 5.2 applies. Additionally, the following holds:*
H: Any $p \in \mathcal{P}$ is connected to exactly two other elements of $\mathcal{P}$. Hence $S \setminus \mathcal{P} = \mathcal{E} \dot\cup \mathcal{J}$, where $\mathcal{E}$ is the exterior of both $\mathcal{P}$ and C, and $\mathcal{J}$ the interior of both $\mathcal{P}$ and C. Moreover, $\mathcal{E}$ is connected, and the same holds for $\mathcal{J}$.

Proof: Since any $p \in \mathcal{P}$ is connected to exactly two other elements of $\mathcal{P}$, then C is thin and spike-free. The same will happen to L, constructed from $\mathcal{P}$ by means of Theorem 6.1. Moreover, L can only be tracked horizontally and vertically. Hence, it is not difficult to conclude that its image $\mathcal{P}'$ has the following property: $S \setminus \mathcal{P}' = \mathcal{E}' \dot\cup \mathcal{J}'$, where $\mathcal{E}'$ is connected, and the same holds for $\mathcal{J}'$. But $\mathcal{E} = \mathcal{E}' \cup \mathcal{L}$, where $\mathcal{L}$ is given by Definition 5.5. Moreover, $\mathcal{E}$ is connected. Now, $\mathcal{J} = S \setminus (\mathcal{P} \cup \mathcal{E}) = S \setminus (\mathcal{P}' \cup \mathcal{E}') = \mathcal{J}'$. Therefore, $\mathcal{J}$ is also connected. q.e.d.

Notice that neither lists of Figure 3 fulfil the hypothesis of Corollary 6.1.

7. Results

This present section is like a manual to the programme. Download `legoloci.zip` from

`https://sites.google.com/site/aefabris/codes`

and extract it in a folder. You will get `loci.p`, `lego.p` and some test-files. In order to draw our figures, we have chosen Xfig for Linux Ubuntu 12.04. They all have extension FIG, and the reader can either change them or even create new tests. However, our programmes *do not* read FIG-files. Whenever you create a figure, please export it in a format like TIF or JPG.

We shall make the programmes `loci.m` and `lego.m` available in future. For the time being, users can run the p-codes in Matlab. To test the robustness of our algorithms, one can either edit the presented test files or even create new ones.

The source codes run both in Matlab and Octave. Table 1 shows the average performance of these programmes in each software, providing you use 4GB of RAM, microprocessor Intel Core i5 3.2GHz, and operational system Linux Ubuntu 12.04. Table 1 refers to Matlab 7.8.0 (R2009a) and Octave 3.2.4 (R2009). All test files have extension JPG, **M** stands for Matlab and **O** for Octave.

Type `loci` at the Matlab prompt. You will be asked to give a filename, for instance `test3.tif`. Please write the full name *with* extension. After pressing the Enter-key, `loci` will display the figure and also its inside according to FUA. The elapsed computational time is printed in the Matlab terminal window. It counts all commands that come right after *having stored* the image in the variable *img* (see Algorithm 1), and right after *having displayed* its interior on the screen.

Now the user gets the following message:

Try CoTRA? Yes = 1; No = any key

Table 1: Comparison to Computational Times in Seconds

	loci-M	**lego-M**	**loci-O**	**lego-O**
test1	0.30	0.80	0.34	2.18
test2	0.36	1.17	0.51	4.03
test3	0.29	0.63	0.34	2.91
test4	0.29	0.76	0.39	3.46
test5	0.56	2.65	1.25	8.40
test6	0.30	0.89	0.41	2.32
test7	0.36	0.59	0.43	7.51
test8	0.30	0.93	0.39	3.86

By choosing 1, CoTRA is called to create the Lego curve as in the proof of Theorem 6.1. A new figure will be displayed, but now its *inside* matches Definition 5.3. Once again, the elapsed computational time is printed in the Matlab terminal window. It counts all commands the same way as done in the previous case.

Now we illustrate some outputs.

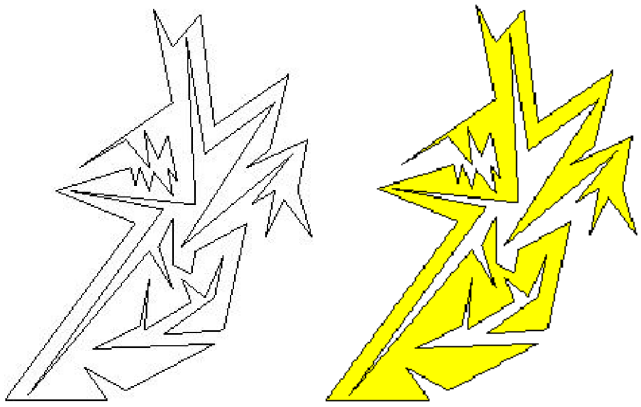

Figure 9: Filling up a challenging picture with CoTRA.

At the Introduction, we mentioned that the marking phase of scan conversion requires careful attention to both geometric and numerical details. Figure 9 left was obtained by drawing a polygonal `test4.fig` in Xfig. Like many other drawing softwares in computer graphics, the continuous mathematical structure is then discretized as a digital picture.

A detail of Figure 9 is zoomed in Figure 6, which shows a self-intersection that occurs when we move from the *ideal simple (i.e. non-self-intersecting) continuous polygon* to the *digital picture*: the digital picture `test4.jpg` is obtained by exporting `test4.fig` to the JPG-format. As commented in Figure 6, CoTRA finds a self-intersection at $(y, x) = (130, 70)$. So `test4.jpg` is filled in accordance with Definition 5.3.

Figure 10 shows that parts with empty or non-connected interiors are correctly detected.

8. Conclusions

In this paper, the difficulties in discretisation are acknowledged and tackled by casting the problem to be solved as a discrete integer problem from the very outset. Integer spaces

Table 2: Comparison to Seed-fill, Scan-conversion and Our approach

	Seed-fill	Scan Conversion	Our approach
Input data	scanned/math defined	math defined	scanned/math defined
Implementation	naïve recursive: simple	complex	simple
Modelling space	screen space	object space	screen space
Boundary modelling	points/polygons	polygons	points
Algorithm	integer/floating point	floating point	integer
Applications areas	image processing CG painting systems	computer graphics	computer graphics and image processing
Additional memory	very large stack	small/large stack	no/small stack
Hard/software design	unsuitable hardware	suitable hard/software	suitable hard/software
Device dependency	requires GetPixelValue	device independent	device independent
Theoretical background	simple	simple	relatively complex
Robustness	to be more robust	most algorithms fail	robust
Efficiency	relatively slow	most are fast	relatively fast

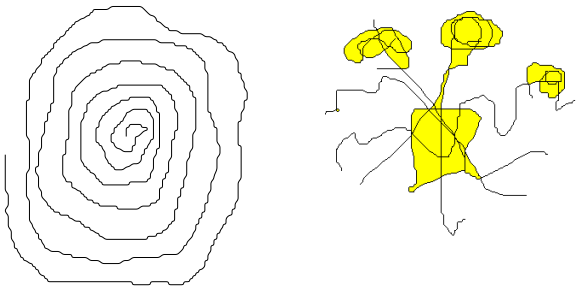

Figure 10: Examples of degenerate regions: empty and non-connected interiors.

are more amenable to analysis and proof. There the commonly employed geometric operations are translated into discrete versions, which allow a better control of robustness.

This approach enables the development of an integer-only algorithm, of which the main characteristics are: simple structure, no special cases to forecast, no extra memory allocation, and proved robustness. Namely, it can handle arbitrarily complex inputs. The source codes `loci.m` and `lego.m` were implemented without either recursive calls or repeat/while-loops. Only for-loops were allowed in the programmes. These characteristics can be forwarded to hardware implementation.

In fact, the main purpose of our work was to make faultless and fast algorithms that could also be implemented in hardware.

Regarding the efficiency of our present implementations, FUA alone has linear complexity with respect to $N \cdot M$. Our observations regarding CoTRA on numerous test scenes suggest quasi-linearity, but before any formalisation we need to check for improvements. Indeed, although CoTRA is relatively fast, we could improve its speed by exploring coherence tests in several ways.

Moreover, FUA and CoTRA require input data that can either be scanned or constructed by a mathematical model. Therefore, these algorithms can be interactively used for applications in both Image Processing and Computer Graphics.

See Table 2 for a summary of seed-fill, scan-conversion and our approach.

Appendix A. Some Theory After This Work

Any seed-fill algorithm easily finds the exterior $\mathcal{E}$ of a picture $\mathcal{P} \subset S$ directly from Definition 5.2. Hence we get the interior as $S \setminus (\mathcal{P} \cup \mathcal{E})$. But the computational cost is huge because seed-fill approaches are recursive.

Theorem 6.1 constructs the Lego curve, which after some pruning is then applied to FUA (see Section 6).

However, the Lego curve is much more than just a way to get correct digital answers. It can help prove theorems in the *continuous* Euclidean plane. For instance, consider this strong version of Jordan's theorem, called *Schöenflies Theorem*:

Theorem Appendix A.1. *Let S^1 be the unitary circumference in $\mathbb{R}^2$ centred at the origin. Let $f : S^1 \to C \subset \mathbb{R}^2$ be a homeomorphism between S^1 and C. Hence, there exists a homeomorphism $F : \mathbb{R}^2 \to \mathbb{R}^2$ such that $F|_{S^1} = f$.*

An immediate consequence of Theorem Appendix A.1 is that the *interior* of C, namely *the bounded component of* $\mathbb{R}^2 \setminus C$, is simply connected. It is a powerful result that, however, does not seem to admit a proof accessible to graduate students. But here we outline some ideas of which details could be carried out in future works.

As depicted in Figure 7, we could get a *continuous* Lego curve out of any grid in which the original C is covered by black pixels. Ideally, resolution becomes infinite when the pixel dimension $(1/n) \times (1/n)$ goes to zero, where n is a positive integer. For each n, we get the continuous Lego curve L_n that corresponds to C. It is not difficult to describe $f_n : S^1 \to L_n$ such that $f_n \to f$, because L_n is tracked only horizontally and vertically. This same reason allows us to construct a homeomorphism $F_n : \mathbb{R}^2 \to \mathbb{R}^2$, with $F_n|_{S^1} = f_n$. Finally, a limit process will then result in the general proof.

References

[1] Schirra S. How reliable are practical point-in-polygon strategies? In: Algorithms-ESA 2008. Springer; 2008, p. 744–55.

[2] Kettner L, Mehlhorn K, Pion S, Schirra S, Yap C. Classroom examples of robustness problems in geometric computations. Computational Geometry 2008;40(1):61 – 78.
[3] Fishkin KP, Barsky BA. A family of new algorithms for soft filling. ACM SIGGRAPH Computer Graphics 1984;18:235–44.
[4] Codrea MC, Nevalainen OS. An algorithm for contour-based region filling. Computer & Graphics 2005;29:441–50.
[5] Tsai YH, Chung KL. Region-filling algorithm on bincode-based contour and its implementation. Computer & Graphics 2000;24(4):529–37.
[6] Lejun S, Hao Z. A new contour fill algorithm for outlined character image generation. Computer & Graphics 1995;19(4):551–6.
[7] Levoy M, Whitted T. The use of points as a display primitive. University of North Carolina, Department of Computer Science; 1985.
[8] Rusinkiewicz S, Levoy M. Qsplat: A multiresolution point rendering system for large meshes. In: Proceedings of the 27th annual conference on Computer graphics and interactive techniques. ACM Press/Addison-Wesley Publishing Co.; 2000, p. 343–52.
[9] Weyrich T, Heinzle S, Aila T, Fasnach DB, Oetiker S, Botscho M, et al. A hardware architecture for surface splatting. ACM Trans Graph 2007;26(3).
[10] Rosenfeld A. Digital topology. American Mathematical Monthly 1979;86(8):621–30.
[11] Kong T, Rosenfeld A. Digital topology: Introduction and survey. Computer Vision, Graphics and Image Processing 1989;48:357–93.